
Rule WLM082: Job might not be suitable for WLM-managed initiators

Finding: CPExpert has detected that a WLM-managed job was spent a major portion of its total active time in a queue waiting for an initiator.

Impact: This finding should be viewed a LOW IMPACT, MEDIUM IMPACT, or HIGH IMPACT on the performance of the batch jobs assigned to the service class. The level of impact will depend primarily on the number of WLM-managed jobs that spent a major portion of their active time in a queue waiting for an initiator.

Logic flow: This is a basic finding. There are no predecessor rules.

Discussion: With OS/390 Version 2 Release 4, the Workload Manager allows installations to define job classes as being managed by Job Entry Subsystem (JES) or by the Workload Manager (WLM).

- For jobs assigned to JES-managed job classes, the normal selection of the jobs for initiation will be done. That is, JES initiators will select jobs from the job class queue based on the normal selection criteria (e.g., priority, aging, resource affinity, etc.).
- For jobs assigned to WLM-managed job classes, the WLM will control selection of jobs by dynamically changing the number of WLM initiators and/or their work selection criteria. These actions will be taken by the WLM in an attempt to meet installation defined goals for the service classes to which the jobs are classified.

WLM control of batch work is enabled by changing the mode of a JES2 job class (on a job class by job class basis) to MODE=WLM. When that is done, no job in that JES2 job class will be selected by normal JES2 initiators. Instead WLM will be informed of the jobs waiting execution and will start and stop WLM controlled initiators based the current backlog of work.

After conversion of a job's JCL, JES2 places the job into an appropriate class queue to await execution. If the job class is a JES2 managed class (JOBCLASS MODE=JES), JES2 initiators select a job from the class queue (depending on the job's priority) and pass control to MVS to execute your program. If the job class is a WLM managed class (JOBCLASS MODE=WLM), JES provides the WLM with a list of jobs waiting to execute, by service class. Based on installation goals for the service classes, the WLM can determine whether job queue time is a significant delay to a

service class meeting its performance goal. If so, the WLM might start an initiator for a job class assigned to the service class.

Job initiators for WLM managed job classes are controlled dynamically by workload management. These initiators run under the Master Subsystem and are not assigned JES2 job numbers. WLM can adjust the number of initiators on each system based on:

- The queue of jobs awaiting execution in WLM managed classes.
- The performance goals and relative importance of this work.
- The success of meeting these goals.
- The capacity of each system to do more work.

The execution queue delay for the jobs will be included in the calculation of execution velocity or response time for the service classes. The resulting execution velocity or response time will be used to calculate a performance index for the service classes. The resulting performance index will be used by the WLM to assess how well the service classes are meeting their performance goals. If a service class is not meeting its performance goal, the WLM might decide to add an initiator to the job class to attempt to reduce the execution queue delay.

When the WLM examines service class delays, it attempts to manage resource allocation to eliminate the most serious delays to work in the service class, where eliminating the delay would cause performance to be significantly improved. It is important to appreciate that for service class periods with execution velocity goals (and for service class periods with response goals under 20 minutes), the WLM will not examine any individual unit of work. Rather, the WLM will examine delays to all work in the service class period¹.

If jobs in a service class spend most of their active time in a queue waiting for initiation, and if this condition is true for most jobs in the service class, the WLM will consider job initiation as the most serious delay for all work in the service class. Consequently, the WLM will tend to ignore other delays (such as CPU delay, MPL delay, etc.) that occur while the jobs are executing. This would have the effect of the WLM not effectively managing the actual execution of the work in the service class!

Prior to OS/390 Version 2 Release 9, only the SMF Type 26 records contained an indication as to whether a job was assigned a job class with

¹There is an exception to this statement in that the working set management algorithms might be applied to a specific unit of work.

MODE=JES or MODE=WLM. Since SMF Type 26 records are not often kept in a performance data base (and are not available until after a job has ended) analysis of conflicts between JES-managed and WLM-managed initiators was not feasible. Consequently, CPExpert requested that IBM place an indicator in the SMF Type 30 records so that potential problems with WLM-managed initiators could be analyzed. With OS/390 Version 2 Release 9, IBM created the SMF30WMI indicator as a part of the SMF30PF1 (performance section flag byte) variable.

CPExpert examines the SMF30WMI indicator in SMF Type 30 records to determine whether a job is assigned to a JES-managed or WLM-managed initiator. CPExpert then computes the percent of the overall active time that the job spent waiting in the execution queue.

CPExpert produces Rule WLM082 if any jobs assigned to a job class with MODE=WLM spent more than 75% of their total active time waiting in the execution queue..

The following example illustrates the output from Rule WLM082:

RULE WLM082: JOBS MIGHT NOT BE SUITABLE FOR WLM-MANAGED INITIATORS							
The execution queue delay for the below jobs accounted for more than 75% percent of the total active time of the jobs in this job class. Since the execution queue delay is so large, the queue delay will dominate the calculation of the Performance Index for the service class. Consequently, the Workload Manager might not address other execution-related delays because they would not significantly affect the performance index. If this finding is regularly produced, you should reevaluate whether this job class should be assigned to WLM-managed initiators or whether the specific jobs should be reassigned to a different job class.							
JOB NAME	JOB CLASS	SYSTEM	SERVICE CLASS	INITIATOR TIME	QUEUE TIME	ACTIVE TIME	
CQQUERY	R	J90	DISCR	17MAR2000:15:01:09	0:00:12	0:00:14	
CQQUERY	R	J90	DISCR	17MAR2000:15:15:29	0:00:12	0:00:14	
CQCHKPT	R	J90	DISCR	17MAR2000:15:15:34	0:00:12	0:00:14	

Suggestion: If Rule WLM082 is regularly produced, CPExpert suggests that you consider the following alternatives:

- If the job active time is small and many other jobs execute in the service class with small execution queue delay times relative to their total active time, you probably should ignore this finding. This means that the execution queue delay time of the specific jobs that were identified will not dominate the Workload Manager's resource management considerations for work in the service class.
- Review the job class limits (the XEQCOUNT=MAXIMUM= parameter on the JOBCLASS statement) to determine whether the job class limits

restrict the number of jobs executing. If the job class limits do restrict the number of jobs executing in the job class, you may wish to increase the job class limit. Of course, other considerations must be reviewed to determine why the job class limit was selected.

- Respecify the job class for the identified job(s), to remove the job(s) from the job class that has MODE=WLM specified.
- Please provide Computer Management Sciences with feed-back on other options that you consider. Thanks!

Reference: OS/390 JES2 Initialization and Tuning Guide
Section 2.1.4.13.3: Considerations for WLM Goals Applied to Batch Jobs

z/OS JES2 Initialization and Tuning Guide |
Section 2.1.4.13.3: Considerations for WLM Goals Applied to Batch Jobs |